

WHAT IS CLAIMED IS:

1 1. A method of detecting surveillance or attack activity over a computer
2 communications network, comprising:

3 receiving a plurality of messages from a data sensor located at a network audit point, each
4 of said messages describing an event occurring on said communications network;
5 classifying one or more of said events to produce one or more labeled alerts;
6 combining one or more said labeled alerts to produce a combined alert; and
7 aggregating one or more said combined alerts to produce an aggregate alert notification.

1 2. The method of claim 1, further comprising filtering one or more said aggregate alert
2 notifications by a cost-based model to produce a qualified alert.

1 3. A method of detecting surveillance activity over a computer communications
2 network, comprising:

3 receiving a plurality of messages from a data sensor located at a network audit point, each
4 of said messages describing an event occurring on said communications network;
5 processing one or more of said messages comprising one or more of the following:
6 clustering packets exchanged between the two addresses within a specified time period;
7 clustering packets exchanged between two addresses having certain flags set;
8 clustering packets exchanged between two addresses having similar flags set; and
9 clustering packets exchanged between two addresses having similar characteristics.

1 4. The method of claim 3, further comprising processing one or more said extrapolated
2 network connections to produce a detected surveillance probe, said processing of one or more
3 said extrapolated network connections to produce a detected surveillance probe comprising
4 one or more of the following:

5 grouping connection session records over related source addresses;
6 scoring each group based on the quantity of attack destinations;
7 generating an alert for each group whose score is greater than an empirically-derived
8 threshold;
9 identifying unusual packets;
10 identifying packets that have a particular arrangement of flags set;
11 identifying packets that have all flags set;
12 identifying packets that have payloads smaller than a predetermined size;
13 identifying packets to which there is no response;
14 identifying packets to which there is no response and that have a particular arrangement
15 of flags set;
16 identifying detected connections with certain characteristics;
17 identifying detected connections with an unusually small number of packets;
18 identifying detected connections with fewer packets than a predetermined limit;
19 identifying detected connections with packets that have traveled only from the source to
20 the destination;
21 identifying detected connections with packets that have traveled only from the destination
22 to the source; and
23 identifying detected connections with packets whose payloads are smaller than a
24 predetermined limit.

1 5. The method of claim 4, further comprising the control of false positive detections vs.
2 false negative detections.

1 6. The method of claim 4, further comprising generation of a profile of surveillance
2 activity, said profile of surveillance activity comprising one or more of the following:

3 a breakdown of probes;
4 the number of attackers;
5 the number of attacks per unit time;
6 the percentage of activity that constitutes malicious surveillance;
7 the breakdown of source country frequencies;
8 the most frequently-targeted network addresses; and
9 the temporal frequency trends of individual attackers.

1 7. The method of claim 4, further comprising processing one or more said detected
2 surveillance probes to produce a detected surveillance scan, said processing of one or more
3 said detected surveillance probes to produce a detected surveillance scan comprising one or
4 more of the following:

5 modeling and detecting surveillance scans as a series of surveillance probes that originate
6 from one or more source addresses and that are sent to one or more destination addresses;

7 modeling and detecting surveillance scans performed by a particular source address by
8 identifying a particular source address that sends more than a specified number of probes;

9 modeling and detecting surveillance scans performed by a particular source address by
10 identifying a source address that generates more than a specified number of probes within
11 a specified time period;

12 modeling and surveillance detecting scans performed by one source IP address by
13 identifying a source address that sends probes to more than a specified number of
14 destinations;

15 modeling and detecting surveillance scans performed by a particular source address by
16 identifying a source address that sends probes to a specified set of destinations;

17 modeling and detecting surveillance scans performed by a particular source address by
18 identifying a source address that sends probes to specified ports;

19 modeling and detecting surveillance scans performed by a particular source address by
20 identifying a source address that sends probes to a number of destinations in excess of a
21 specified limit within a specified time period;

22 limiting the number of detected scans by reporting only source addresses that perform
23 more than a specified number of probes within a specified time; and

24 limiting the number of detected scans by reporting only source address groups that
25 perform more than a specified number of probes within a specified time.

1 8. The method of claim 7, further comprising the control of false positive detections vs.
2 false negative detections.

1 9. The method of claim 7, further comprising generation of a profile of surveillance
2 activity, said profile of surveillance activity comprising one or more of the following:
3 a breakdown of probes;
4 a breakdown of scans;
5 the number of attackers;
6 the number of attacks per unit time;
7 the percentage of activity that constitutes malicious surveillance;
8 the breakdown of source country frequencies;
9 the most frequently-targeted network addresses; and
10 the temporal frequency trends of individual attackers.

1 10. The method of claim 7, further comprising processing one or more said detected
2 surveillance scans to detect a group of scanning hosts, said processing of one or more said
3 detected surveillance scans to detect a group of scanning hosts comprising:
4 modeling and detecting scans distributed across a series of source addresses by grouping
5 addresses, said grouping of addresses being performed by subtracting one address from
6 another and placing the two addresses in the same group if the difference is less than a
7 specified amount.

1 11. The method of claim 10, further comprising the control of false positive detections vs.
2 false negative detections.

1 12. The method of claim 10, further comprising generation of a profile of surveillance
2 activity, said profile of surveillance activity comprising one or more of the following:

3 a breakdown of probes;
4 a breakdown of scans;
5 the number of attackers;
6 the number of attacks per unit time;
7 the percentage of activity that constitutes malicious surveillance;
8 the breakdown of source country frequencies;
9 the most frequently-targeted network addresses; and
10 the temporal frequency trends of individual attackers.

1 13. A method of detecting surveillance or attack activity over a communication network
2 comprising:

3 combining alerts to such surveillance or attack activity generated by an intrusion
4 detection system with alerts to such surveillance or attack activity generated by an
5 anomaly detection system to produce a combined alert;
6 prioritizing said combined alert to produce a prioritized alert;
7 presenting said prioritized alert to a security analyst.

1 14. A computer program product for use in conjunction with a computer system to
2 classify and analyze surveillance or attack activity over a communications network, the
3 computer program product comprising a computer readable storage medium and a computer
4 program mechanism embedded therein, the computer program mechanism comprising:

5 an event data storage buffer that receives and stores incoming event data;
6 an initial event evaluator that receives event data from said event data storage buffer and
7 generates raw alerts;
8 a raw alert data storage buffer that receives and stores said raw alerts;
9 a post-processing alert evaluator that receives said stored raw alerts and produces
10 processed alerts;
11 a plurality of alert filtering modules that receive said processed alerts and produce user
12 alerts;
13 a user alert data buffer that receives and stores said user alerts;

14 a plurality of production models for said initial event evaluator;
15 a plurality of production models for said alert filtering modules;
16 storage for said production models for said initial event evaluator and for said production
17 models for said alert filtering modules; and
18 an automated job submission manager that orchestrates the operations of said initial event
19 evaluator and of said post-processing alert evaluator.

1 15. A computer system for formatting, classifying and analyzing surveillance or attacks
2 over a communications network, the computer system comprising:
3 a central processing unit;
4 a memory, coupled to the central processing unit, the memory storing:
5 outputs of sensors connected to the communications network;
6 outputs of an initial event evaluator;
7 outputs of a post-processing alert evaluator;
8 outputs of a plurality of alert filtering modules;
9 a plurality of production models for said initial event evaluator; and
10 a plurality of production models for said alert filtering modules.

1 16. A method of processing computer network surveillance alerts, comprising:
2 receiving alerts from an intrusion detection system;
3 receiving alerts from an anomaly detection system;
4 receiving alerts from a scan / probe detection system;
5 aggregating one or more of said alerts from said intrusion detection system, said anomaly
6 detection system, and said scan / probe detection system; and
7 generating an aggregated alert.

1 17. A user display for profiling surveillance activity over a computer network, said user
2 display comprising: a display of a numerical estimate of the severity of an attack and one or
3 more of the following:
4 a list of the highest priority threats;

- 5 a list of the highest priority targets;
- 6 detailed threat information;
- 7 detailed target information;
- 8 the country of origin of an attack;
- 9 the country of origin of a target; and
- 10 a plot of attack severity versus time.

- 1 18. A method of detecting surveillance or attack activity over a computer
- 2 communications network, comprising:
- 3 modeling network connections;
- 4 detecting said network connections that are likely surveillance probes originating from
- 5 malicious sources;
- 6 detecting scanning activity by grouping source addresses that are logically close to one
- 7 another; and
- 8 recognizing certain combinations of said likely surveillance probes.